Conditionals: between language and reasoning

Class 9: background semantics

June 18, 2019

Breaking de Morgan's law in antecedents

Inquisitiveness and counterfactuals

- An antecedent is associated with a set of alternatives.
- Each alternative counts as a separate counterfactual assumption.
- The counterfactual is true if the consequent follows on each assumption.

Inquisitiveness and counterfactuals

- An antecedent is associated with a set of alternatives.
- Each alternative counts as a separate counterfactual assumption.
- The counterfactual is true if the consequent follows on each assumption.

Implementation:

Let $f(\cdot,\cdot)$ be a selection function.

$$w \Vdash \varphi > \psi \iff \forall a \in \mathrm{Alt}(\varphi) : f(w,a) \subseteq |\psi|$$

Back to our switches

Here are the inquisitive semantics representations for our antecedents:

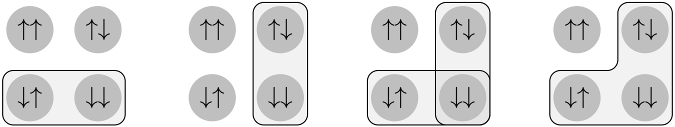

(a) $\overline{A}$ (b) $\overline{B}$ (c) $\overline{A} \vee \overline{B}$ (d) $\neg(A \wedge B)$

Back to our switches

Here are the inquisitive semantics representations for our antecedents:

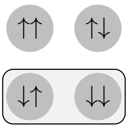

(e) $\overline{A}$

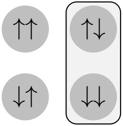

(f) $\overline{B}$

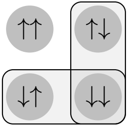

(g) $\overline{A} \vee \overline{B}$

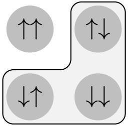

(h) $\neg(A \wedge B)$

- This theory breaks de Morgan's law: $\overline{A} \vee \overline{B}$ and $\neg(A \wedge B)$ have the same truth-conditions, but different alternatives.

Back to our switches

Here are the inquisitive semantics representations for our antecedents:

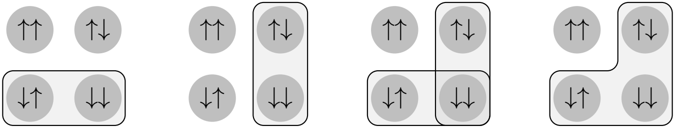

(i) $\overline{A}$ (j) $\overline{B}$ (k) $\overline{A} \vee \overline{B}$ (l) $\neg(A \wedge B)$

- This theory breaks de Morgan's law: $\overline{A} \vee \overline{B}$ and $\neg(A \wedge B)$ have the same truth-conditions, but different alternatives.
- The antecedents $\overline{A}$, $\overline{B}$, and $\neg(A \wedge B)$ are non-inquisitive: so the corresponding counterfactuals are interpreted standardly.

Back to our switches

Here are the inquisitive semantics representations for our antecedents:

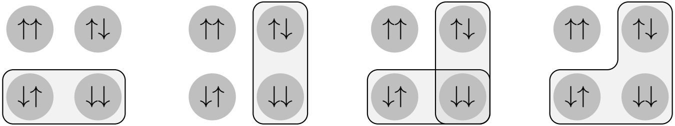

(m) $\overline{A}$ (n) $\overline{B}$ (o) $\overline{A} \vee \overline{B}$ (p) $\neg(A \wedge B)$

- This theory breaks de Morgan's law: $\overline{A} \vee \overline{B}$ and $\neg(A \wedge B)$ have the same truth-conditions, but different alternatives.
- The antecedents $\overline{A}$, $\overline{B}$, and $\neg(A \wedge B)$ are non-inquisitive: so the corresponding counterfactuals are interpreted standardly.
- By contrast, $\overline{A} \vee \overline{B}$ is inquisitive: inquisitive lifting will ensure that $\overline{A} \vee \overline{B} > \mathit{Off}$ is interpreted as equivalent to $(\overline{A} > \mathit{Off}) \wedge (\overline{B} > \mathit{Off})$.

- Now suppose our selection function f is such that:
  - $f(w, |\overline{A}|) \subseteq |Off|$
  - $f(w, |\overline{B}|) \subseteq |Off|$
  - $f(w, |\neg(A \wedge B)|) \not\subseteq |Off|$

- The selection function approach gives the right verdicts for (1), (2), (4):

(1)	$\overline{A} > Off$	✓
(2)	$\overline{B} > Off$	✓
(3)	$\overline{A} \vee \overline{B} > Off$	✓
(4)	$\neg(A \wedge B) > Off$	✗

- Now suppose our selection function f is such that:
  - $f(w, |\overline{A}|) \subseteq |Off|$
  - $f(w, |\overline{B}|) \subseteq |Off|$
  - $f(w, |\neg(A \wedge B)|) \not\subseteq |Off|$
- The selection function approach gives the right verdicts for (1), (2), (4):

(1)	$\overline{A} > Off$	✓
(2)	$\overline{B} > Off$	✓
(3)	$\overline{A} \vee \overline{B} > Off$	✓
(4)	$\neg(A \wedge B) > Off$	✗

- The inquisitive account will make the same predictions for these sentences.
- But the original account would predict (3) to be false, since (3) $\equiv$ (4); the inquisitive account will predict (3) to be true, since (3) $\equiv$ (1) $\wedge$ (2).

Theorem

No minimal change selection function validates:

- $f(w, |\overline{A}|) \subseteq |Off|$
- $f(w, |\overline{B}|) \subseteq |Off|$
- $f(w, |\neg(A \wedge B)|) \not\subseteq |Off|$

Thus, making assumption does not work by maximizing "similarity", regardless of how exactly similarity is construed.

Theorem
No minimal change selection function validates:

- $f(w, |\overline{A}|) \subseteq |Off|$
- $f(w, |\overline{B}|) \subseteq |Off|$
- $f(w, |\neg(A \wedge B)|) \not\subseteq |Off|$

Thus, making assumption does not work by maximizing "similarity", regardless of how exactly similarity is construed.

Question
What procedure do we follow in making counterfactual assumptions? In other words, how to define the relevant selection function $f(\cdot, \cdot)$?

Background semantics

If switch A was down...

- we hold the position of B fixed.

If switch A and switch B were not both up...

- we consider all positions for the switches,
  not only those most similar to the actual one.

If switch A was down...

- we hold the position of B fixed.

If switch A and switch B were not both up...

- we consider all positions for the switches, not only those most similar to the actual one.

So...

- There seems to be no general pressure to consider only scenarios that are minimally different from the actual world.

If switch A was down...

- we hold the position of B fixed.

If switch A and switch B were not both up...

- we consider all positions for the switches,
  not only those most similar to the actual one.

So...

- There seems to be no general pressure to consider only scenarios that are minimally different from the actual world.
- But then, why do we hold the position of B fixed when faced with the assumption that A was down?

If switch A was down. . .

- we hold the position of B fixed.

If switch A and switch B were not both up. . .

- we consider all positions for the switches,
  not only those most similar to the actual one.

So. . .

- There seems to be no general pressure to consider only scenarios that are minimally different from the actual world.
- But then, why do we hold the position of B fixed when faced with the assumption that A was down?
- Proposal: in that case, the position of B is not called into question. It can be regarded as part of the background for the assumption.

Background semantics

- In background semantics we abandon the "minimal change principle", proposing instead a distinction between background and foreground facts.
- When faced with an assumption, we determine which facts to manipulate (foreground) and which to leave alone (background).
- Background facts are held fixed; foreground facts are allowed to change, and their change is not subject to any minimality constraint.

Background semantics

- In background semantics we abandon the "minimal change principle", proposing instead a distinction between background and foreground facts.
- When faced with an assumption, we determine which facts to manipulate (foreground) and which to leave alone (background).
- Background facts are held fixed; foreground facts are allowed to change, and their change is not subject to any minimality constraint.
- $\varphi > \psi$ is true iff ψ follows by causal laws from φ + background facts: this can be seen as an implementation of the meta-linguistic theory.

Principles ruling the background/foreground split:

Principles ruling the background/foreground split:

1. Facts responsible for the falsity of the assumption must be foregrounded.
   - Give up at least those facts which are in conflict with the assumption.

Principles ruling the background/foreground split:

1. Facts responsible for the falsity of the assumption must be foregrounded.
   - Give up at least those facts which are in conflict with the assumption.
2. Causal consequences of foregrounded facts must be foregrounded.
   - If we give up a fact, we must also give up things that depend on it.

Principles ruling the background/foreground split:

1. Facts responsible for the falsity of the assumption must be foregrounded.
   - Give up at least those facts which are in conflict with the assumption.
2. Causal consequences of foregrounded facts must be foregrounded.
   - If we give up a fact, we must also give up things that depend on it.
3. By default, facts are backgrounded.
   - Don't give up facts without a reason.

Causal models (drawing on Pearl 00, Kaufmann 13)
A causal model has a set V of causal variables and a set L of causal laws.

Causal variables
A causal variable X is a partition of the logical space.

- The cells are called settings.
- The true setting of X at w is called the value of X at w, denoted X_w.

In our context, $V = \{?A, ?B, ?On\}$:

- $?A = \{A, \overline{A}\}$
- $?B = \{B, \overline{B}\}$
- $?On = \{On, Off\}$

Causal laws

A causal law l is a triple $\langle C_l, E_l, m_l \rangle$, where:

- C_l is a set of variables (the causes)
- E_l is a variable (the effect)
- m_l is a map from settings of C_l to settings of E_l

In our context, there is only one law.

- Causes: $?A, ?B$
- Effect: $?On$
- Map:

$$\begin{array}{ll} A, B \mapsto On & A, \overline{B} \mapsto Off \\ \overline{A}, B \mapsto Off & \overline{A}, \overline{B} \mapsto On \end{array}$$

Causal graph:

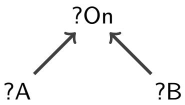

Facts

The facts at w are values of the causal variables:

$$\mathcal{F}_w = \{X_w \mid X \in V\}$$

In our scenario:

$\mathcal{F}_w = \{A, B, On\}$

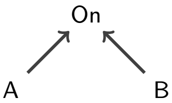

Fact f contributes to the falsity of assumption a at world w if there exists some set F' of facts such that:

- F' is logically consistent with a
- $F' \cup \{f\}$ is logically inconsistent with a

Fact f contributes to the falsity of assumption a at world w if there exists some set F' of facts such that:

- F' is logically consistent with a
- $F' \cup \{f\}$ is logically inconsistent with a

Assump.	Facts contributing to its falsity
$\overline{A}$	A
$\overline{B}$	B
$\neg(A \wedge B)$	A, B

Background conditions

A set $\mathcal{B}(w, a) \subseteq \mathcal{F}_w$ is a background for assumption a at w if it satisfies:

1. Facts responsible for falsity of the antecedent are foregrounded:
   if f contributes to the falsity of a then $f \notin \mathcal{B}(w, a)$
2. Consequences of foregrounded facts are foregrounded:
   if $f \notin \mathcal{B}(w, a)$ and f' is causally dependent on f, then $f' \notin \mathcal{B}(w, a)$.

Background conditions

A set $\mathcal{B}(w, a) \subseteq \mathcal{F}_w$ is a background for assumption a at w if it satisfies:

1. Facts responsible for falsity of the antecedent are foregrounded:
   if f contributes to the falsity of a then $f \notin \mathcal{B}(w, a)$

2. Consequences of foregrounded facts are foregrounded:
   if $f \notin \mathcal{B}(w, a)$ and f' is causally dependent on f, then $f' \notin \mathcal{B}(w, a)$.

For any assumption there is a greatest background:

- $\mathcal{B}^{max}(w, a) \;=\; \mathcal{F}_w - \{f \mid f$ contributes to the falsity of a or f is dependent on such a fact$\}$

Background conditions

A set $\mathcal{B}(w, a) \subseteq \mathcal{F}_w$ is a background for assumption a at w if it satisfies:

1. Facts responsible for falsity of the antecedent are foregrounded:
   if f contributes to the falsity of a then $f \notin \mathcal{B}(w, a)$

2. Consequences of foregrounded facts are foregrounded:
   if $f \notin \mathcal{B}(w, a)$ and f' is causally dependent on f, then $f' \notin \mathcal{B}(w, a)$.

For any assumption there is a greatest background:

- $\mathcal{B}^{max}(w, a) = \mathcal{F}_w - \{f \mid f$ contributes to the falsity of a or f is dependent on such a fact$\}$

Assump.	$\mathcal{B}^{max}(w, a)$
$\overline{A}$	$\{B\}$
$\overline{B}$	$\{A\}$
$\neg(A \wedge B)$	$\emptyset$

Hypothetical context

Making an assumption a in world w given a background $\mathcal{B}(w, a)$ results in the hypothetical context $f_{\mathcal{B}}(w, a)$ consisting of worlds where:

- a is true
- all facts in $\mathcal{B}(w, a)$ are true
- the causal laws are obeyed (simplification)

Hypothetical context
Making an assumption a in world w given a background $\mathcal{B}(w, a)$ results in the hypothetical context $f_{\mathcal{B}}(w, a)$ consisting of worlds where:

- a is true
- all facts in $\mathcal{B}(w, a)$ are true
- the causal laws are obeyed (simplification)

So, now we have a theory of how the selection function operates. We can then plug in this theory into the inquisitive lifting recipe.

Assessment of a counterfactual

Given a causal model M, a background map $\mathcal{B}$, and a world w:

$$M, w \models \varphi > \psi \quad \Longleftrightarrow \quad \forall a \in \mathsf{Alt}(\varphi) : f_{\mathcal{B}}(w, a) \subseteq |\psi|$$

Assessment of a counterfactual

Given a causal model M, a background map $\mathcal{B}$, and a world w:

$$M, w \models \varphi > \psi \quad \iff \quad \forall a \in \mathrm{Alt}(\varphi) : f_{\mathcal{B}}(w, a) \subseteq |\psi|$$

NB

Notice that here φ and ψ are arbitrary: unlike Pearl's, this account is not restricted to special antecedents.

Alternatives for our antecedents, in InqSem

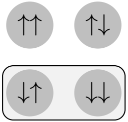

(q) $\overline{A}$

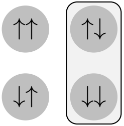

(r) $\overline{B}$

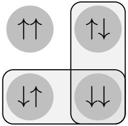

(s) $\overline{A} \vee \overline{B}$

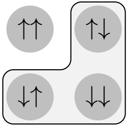

(t) $\neg(A \wedge B)$

$\overline{A} > Off$

$\overline{A} > Off$

- Assumption: $\overline{A}$

$\overline{A} > Off$

- Assumption: $\overline{A}$
- Maximal background: $\{B\}$

$\overline{A} > Off$

- Assumption: $\overline{A}$
- Maximal background: $\{B\}$
- Hypothetical context: $|\overline{A}| \cap |B| \cap |law| \subseteq |Off|$

$\overline{A} > Off$

- Assumption: $\overline{A}$
- Maximal background: $\{B\}$
- Hypothetical context: $|\overline{A}| \cap |B| \cap |law| \subseteq |Off|$
- True.

$\overline{A} > \mathit{Off}$

- Assumption: $\overline{A}$
- Maximal background: $\{B\}$
- Hypothetical context: $|\overline{A}| \cap |B| \cap |\mathit{law}| \subseteq |\mathit{Off}|$
- True.

$\overline{B} > \mathit{Off}$

- True, analogously.

$\overline{A} > Off$

- Assumption: $\overline{A}$
- Maximal background: $\{B\}$
- Hypothetical context: $|\overline{A}| \cap |B| \cap |law| \subseteq |Off|$
- True.

$\overline{B} > Off$

- True, analogously.

$\overline{A} \vee \overline{B} > Off$

- Assumptions: $\overline{A}$, $\overline{B}$

$\overline{A} > Off$

- Assumption: $\overline{A}$
- Maximal background: $\{B\}$
- Hypothetical context: $|\overline{A}| \cap |B| \cap |law| \subseteq |Off|$
- True.

$\overline{B} > Off$

- True, analogously.

$\overline{A} \vee \overline{B} > Off$

- Assumptions: $\overline{A}$, $\overline{B}$
- Under both assumptions, the consequent follows.

$\overline{A} > Off$

- Assumption: $\overline{A}$
- Maximal background: $\{B\}$
- Hypothetical context: $|\overline{A}| \cap |B| \cap |law| \subseteq |Off|$
- True.

$\overline{B} > Off$

- True, analogously.

$\overline{A} \vee \overline{B} > Off$

- Assumptions: $\overline{A}$, $\overline{B}$
- Under both assumptions, the consequent follows.
- True.

$\neg(A \wedge B) > \mathit{Off}$

$\neg(A \wedge B) > Off$

- Assumption: $\neg(A \wedge B)$

$\neg(A \wedge B) > Off$

- Assumption: $\neg(A \wedge B)$
- Maximal background: $\emptyset$

$\neg(A \wedge B) > \mathit{Off}$

- Assumption: $\neg(A \wedge B)$
- Maximal background: $\emptyset$
- Hypothetical context: $|\neg(A \wedge B)| \cap |\mathit{law}| \not\subseteq |\mathit{Off}|$

$\neg(A \wedge B) > Off$

- Assumption: $\neg(A \wedge B)$
- Maximal background: $\emptyset$
- Hypothetical context: $|\neg(A \wedge B)| \cap |law| \not\subseteq |Off|$
- Not true.

Predictions: background theory + inquisitive lifting

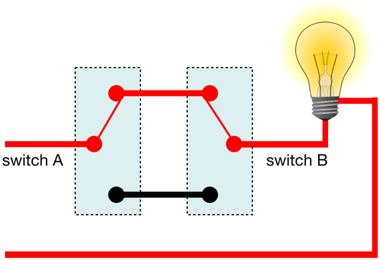

(1) If switch A was down, the light would be off. ✓

(2) If switch B was down, the light would be off. ✓

(3) If switch A or switch B was down, the light would be off. ✓

(4) If switch A and switch B were not both up, the light would be off. ✗

More evidence for the background/foreground divide

More evidence for the background/foreground divide

(5) If I suddenly became taller than I am, you would not notice.

More evidence for the background/foreground divide

(5) If I suddenly became taller than I am, you would not notice. ?

More evidence for the background/foreground divide

(5) If I suddenly became taller than I am, you would not notice. ?

(6) If I suddenly became much taller than I am, you would notice.

More evidence for the background/foreground divide

(5) If I suddenly became taller than I am, you would not notice. ?

(6) If I suddenly became much taller than I am, you would notice. ✓

More evidence for the background/foreground divide

(5) If I suddenly became taller than I am, you would not notice. ?

(6) If I suddenly became much taller than I am, you would notice. ✓

- (5) shows that in changing my height we don't minimize difference.

More evidence for the background/foreground divide

(5) If I suddenly became taller than I am, you would not notice. ?

(6) If I suddenly became much taller than I am, you would notice. ✓

- (5) shows that in changing my height we don't minimize difference.
- (6) shows that we hold fixed your ability to pick up differences in height.

More evidence for the background/foreground divide

(5) If I suddenly became taller than I am, you would not notice. ?

(6) If I suddenly became much taller than I am, you would notice. ✓

- (5) shows that in changing my height we don't minimize difference.
- (6) shows that we hold fixed your ability to pick up differences in height.
- To make the right predictions in minimal change semantics we can stipulate that your perception matters for similarity, but not my height. But why would that be?

More evidence for the background/foreground divide

(5)	If I suddenly became taller than I am, you would not notice.	?
(6)	If I suddenly became much taller than I am, you would notice.	✓

- (5) shows that in changing my height we don't minimize difference.
- (6) shows that we hold fixed your ability to pick up differences in height.
- To make the right predictions in minimal change semantics we can stipulate that your perception matters for similarity, but not my height. But why would that be?
- Background semantics provides a better explanation:
  - my height is a foregrounded variable, and can change freely;
  - your perceptual abilities are backgrounded and thus held fixed.

Evidence for non-maximal background

- In our main experiment, each participant saw one of the target sentences, as well as the filler, (7):

(7) If switch A and switch B were both down, the light would be off.

$\overline{A} \wedge \overline{B} > \mathit{Off}$

- Some participants saw the target item first, while others saw the filler first.

Ordering effects

Table 7: Order effects in the main experiment: target precedes filler

Sentence	Number	True	(%)	False	(%)	Indet.	(%)
$\overline{A} >$ OFF	125	100	80%	3	2.4%	22	17.6%
$\overline{B} >$ OFF	124	94	75.81%	4	3.22%	26	20.97%
$\overline{A} \vee \overline{B} >$ OFF	185	146	78.92%	9	4.86%	30	16.22%
$\neg(A \wedge B) >$ OFF	193	38	19.69%	82	42.49%	73	37.82%
$\neg(A \wedge B) >$ ON	102	21	20.59%	35	34.31%	46	45.10%

Table 8: Order effects in the main experiment: filler precedes target

Sentence	Number	True	(%)	False	(%)	Indet.	(%)
$\overline{A} >$ OFF	131	69	52.67%	3	2.29%	59	45.04%
$\overline{B} >$ OFF	111	59	53.15%	3	2.70%	49	44.14%
$\overline{A} \vee \overline{B} >$ OFF	177	105	59.32%	5	2.82%	67	37.85%
$\neg(A \wedge B) >$ OFF	179	44	24.58%	54	30.17%	81	45.25%
$\neg(A \wedge B) >$ ON	98	22	22.45%	28	28.57%	48	48.98%

- The order effects can be explained assuming that, in addition to the antecedent, other factors may lead to foregrounding certain facts.
- In our case, the filler is: $\overline{A} \wedge \overline{B} > \mathit{Off}$:

(8) If switch A and switch B were both down, the light would be off.

- The assumption $\overline{A} \wedge \overline{B}$ foregrounds the position of both switches.
- To some participants, once the position of B has been foregrounded, it remains foregrounded when interpreting the antecedent $\overline{A}$.
- This leads to empty background, and to judge $\overline{A} > \mathit{Off}$ as 'indeterminate'.
- The same explanation works for $\overline{B} > \mathit{Off}$ and for $\overline{A} \vee \overline{B} > \mathit{Off}$.

- Why no order effects for $\neg(A \wedge B) > \mathit{Off}$ and $\neg(A \wedge B) > \mathit{On}$?
- The assumption $\neg(A \wedge B)$ already foregrounds all the facts.
- Interpreting the filler cannot lead to more facts being foregrounded.

- Why no order effects for $\neg(A \wedge B) > \mathit{Off}$ and $\neg(A \wedge B) > \mathit{On}$?
- The assumption $\neg(A \wedge B)$ already foregrounds all the facts.
- Interpreting the filler cannot lead to more facts being foregrounded.
- These data indicate that in order to decide what to background, we take into account not just the antecedent, but also the preceding discourse.
(cf. the dynamic accounts of von Fintel 2001, Gillies 2007)